

SC2000 Network Challenge

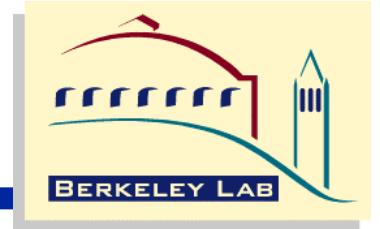
Winner: Fastest and Fattest Award

Brian Tierney, Wes Bethel, Jason Lee, Dan Gunter
Lawrence Berkeley National Laboratory

Helen Chen, Jim Brandt
Sandia National Laboratory

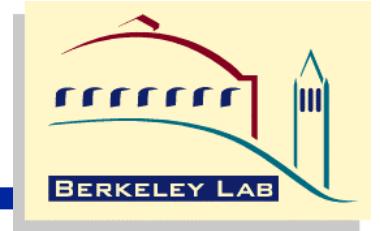
George Pavel, Brian Bodtker
Lawrence Livermore National Laboratory

Challenge Application



- Server (located at LBNL)
 - Distributed Parallel Storage System (DPSS)
- Clients (located on SC2000 show floor)
 - Visapult: Remote Visualization of an 80 GB data set
 - running on an 8 CPU SGI Origin, ACSI booth,
 - <http://www-vis.lbl.gov/projects/visapult/index.html>
 - dpss_get: high-speed parallel file transfer application
 - running on 8 node Linux cluster, ANL booth
- Results:
 - Transferred 262 GB of data in 60 minutes from LBL to SC2000 show floor (Dallas, TX)
 - 5 sec peak throughput: 1.48 Gbits/sec (user data)
 - 60 minute average throughput: 582 Mbits/sec

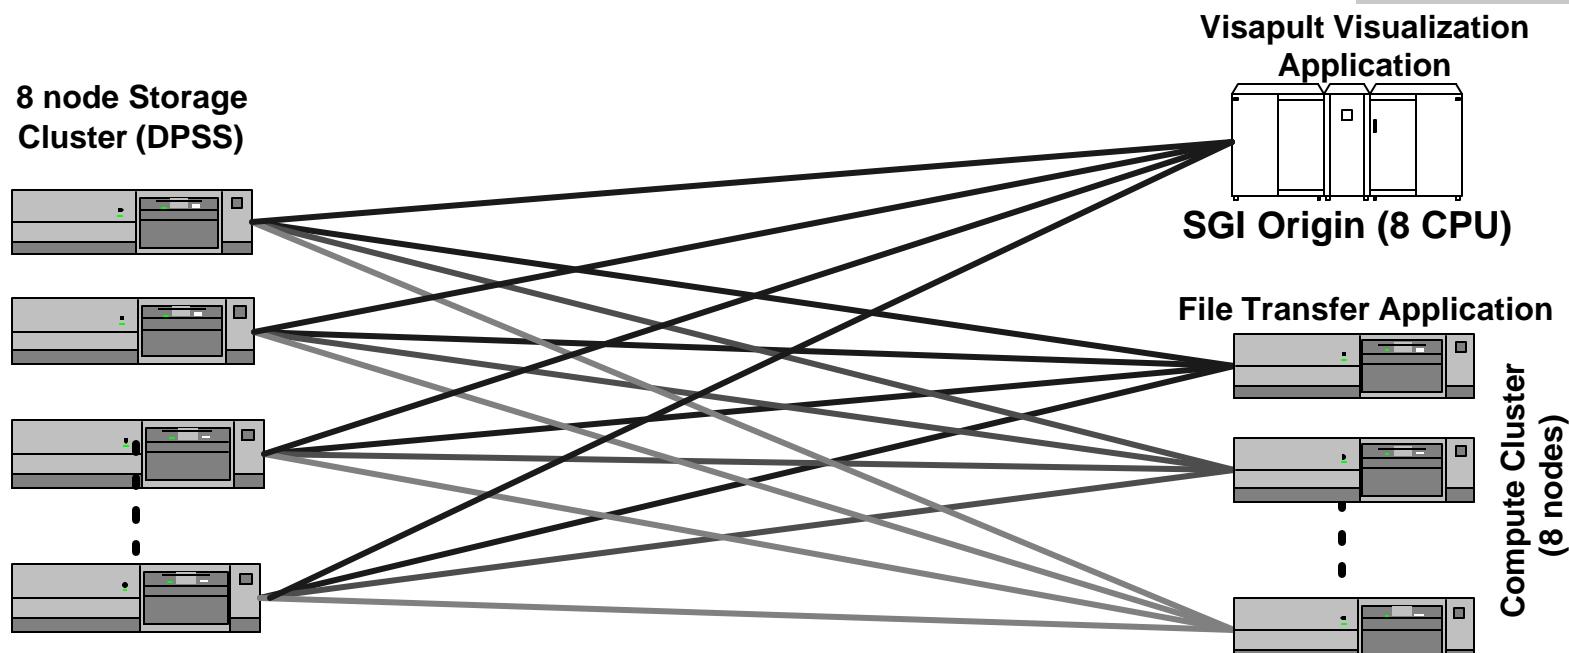
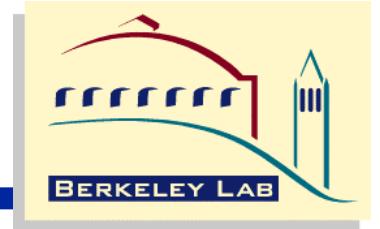
Challenge Application



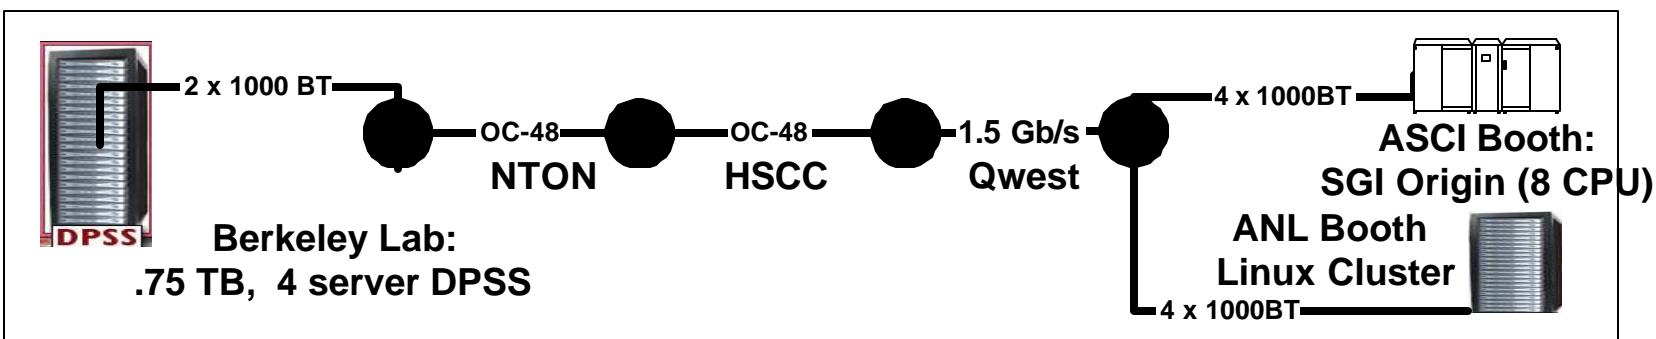
- Ran Visapult and multiple dpss_get processes simultaneously
 - I/O limit on Octane running Visapult about 700-800 Mbits/sec
 - Used dpss_gets to utilize remaining DPSS and network capacity

SC2000 Network Challenge

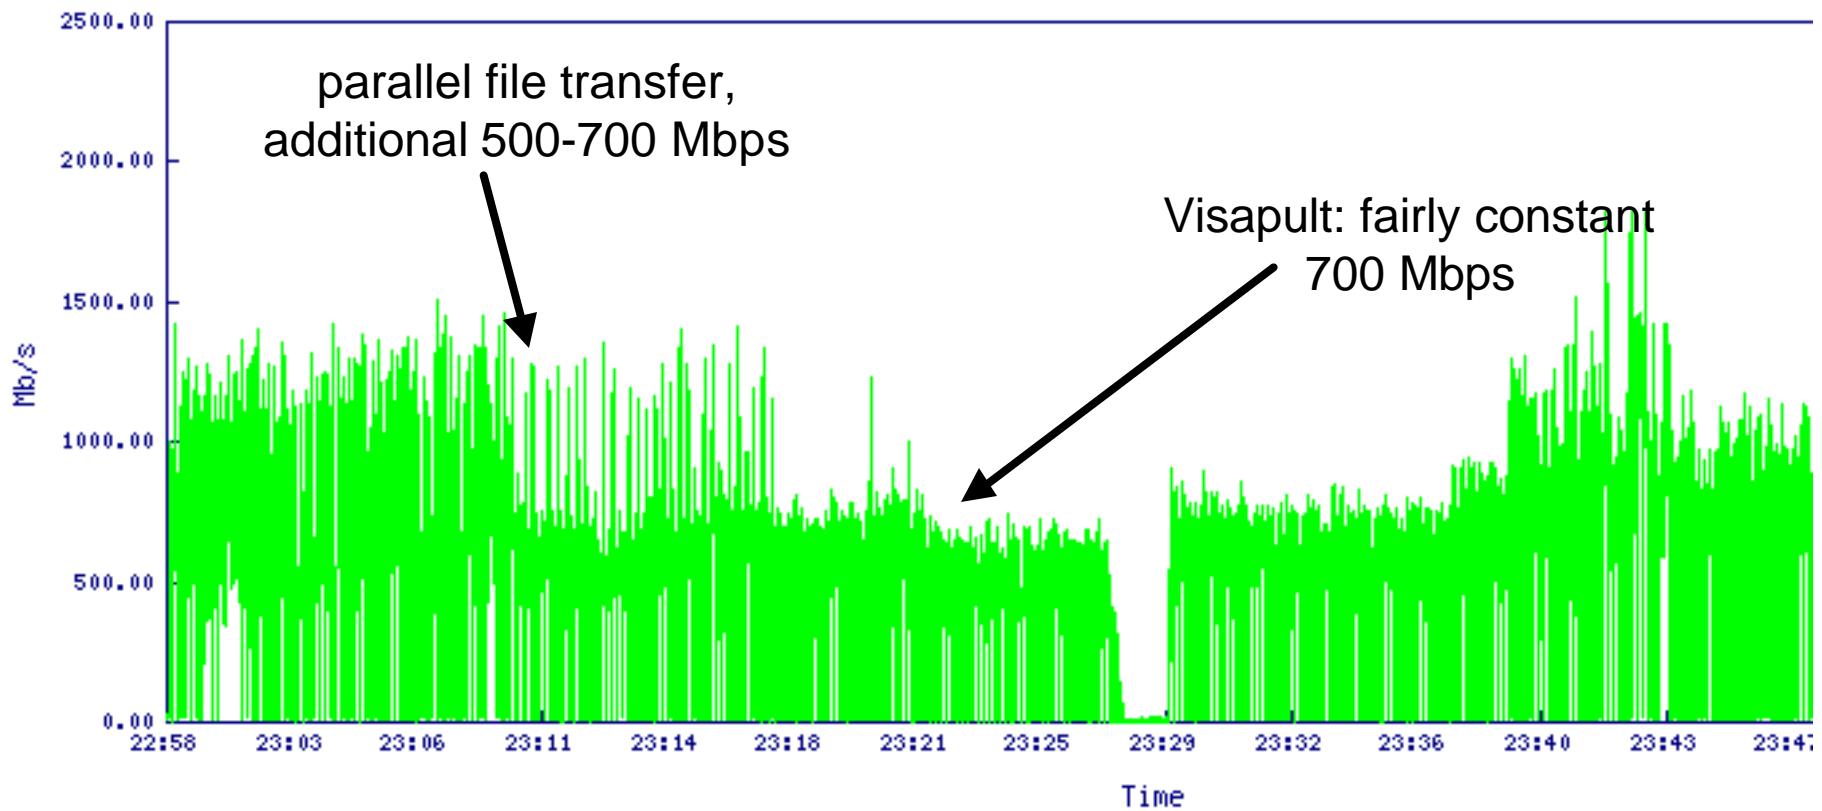
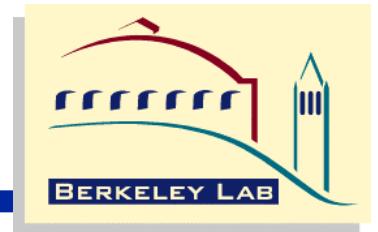
Config: Peak BW = 1.48 Gbits/sec



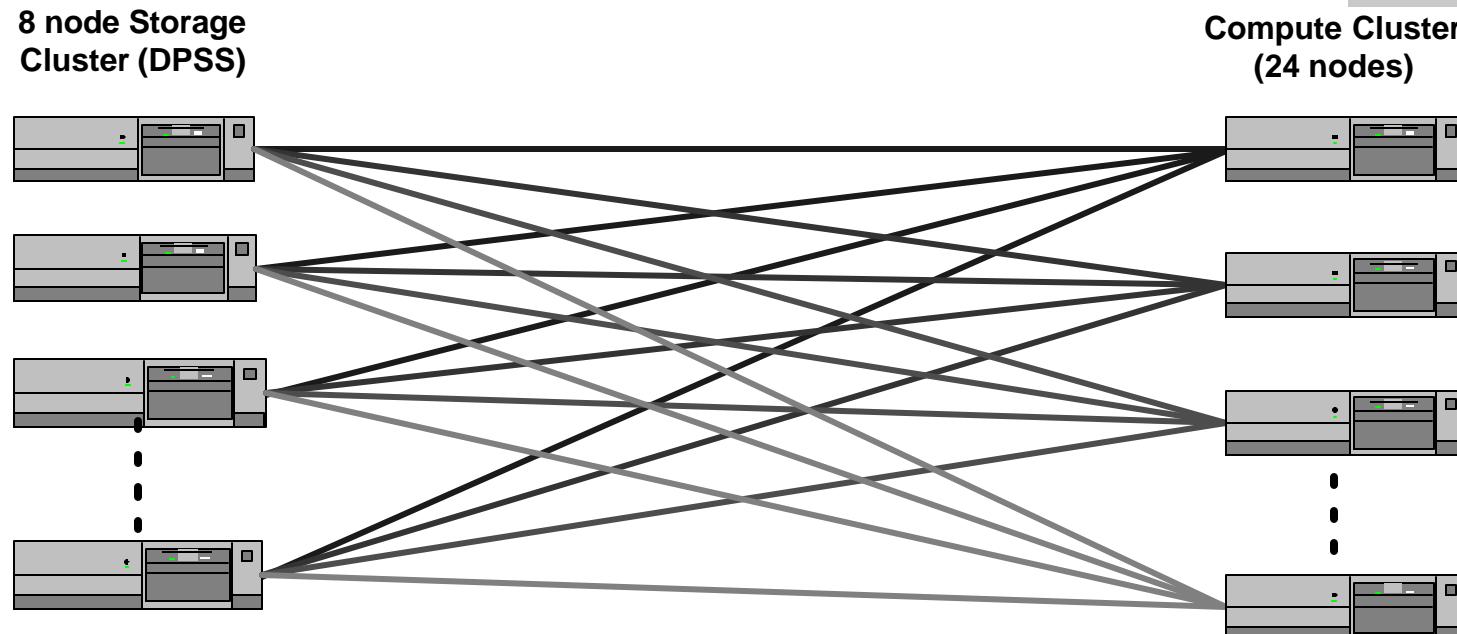
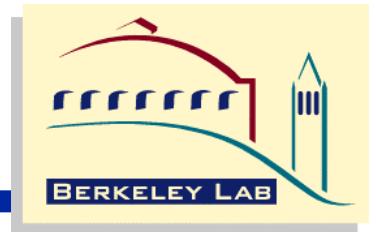
Network Throughput: 5 sec peak 1.48 Gbits/sec (72 streams: 20.5 Mbits/stream); 60 minute sustained average: 582 Mbits/sec



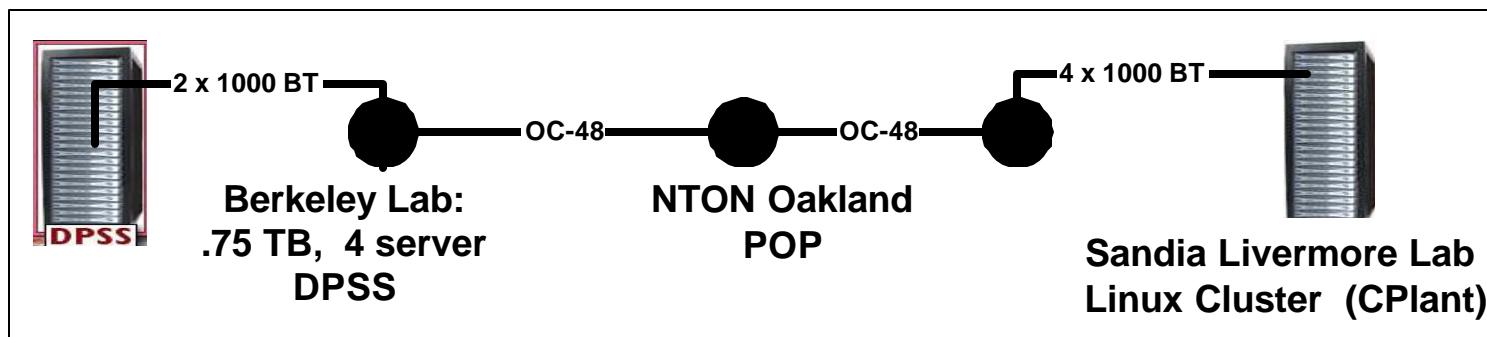
Bandwidth Statistics reported by SciNet Router



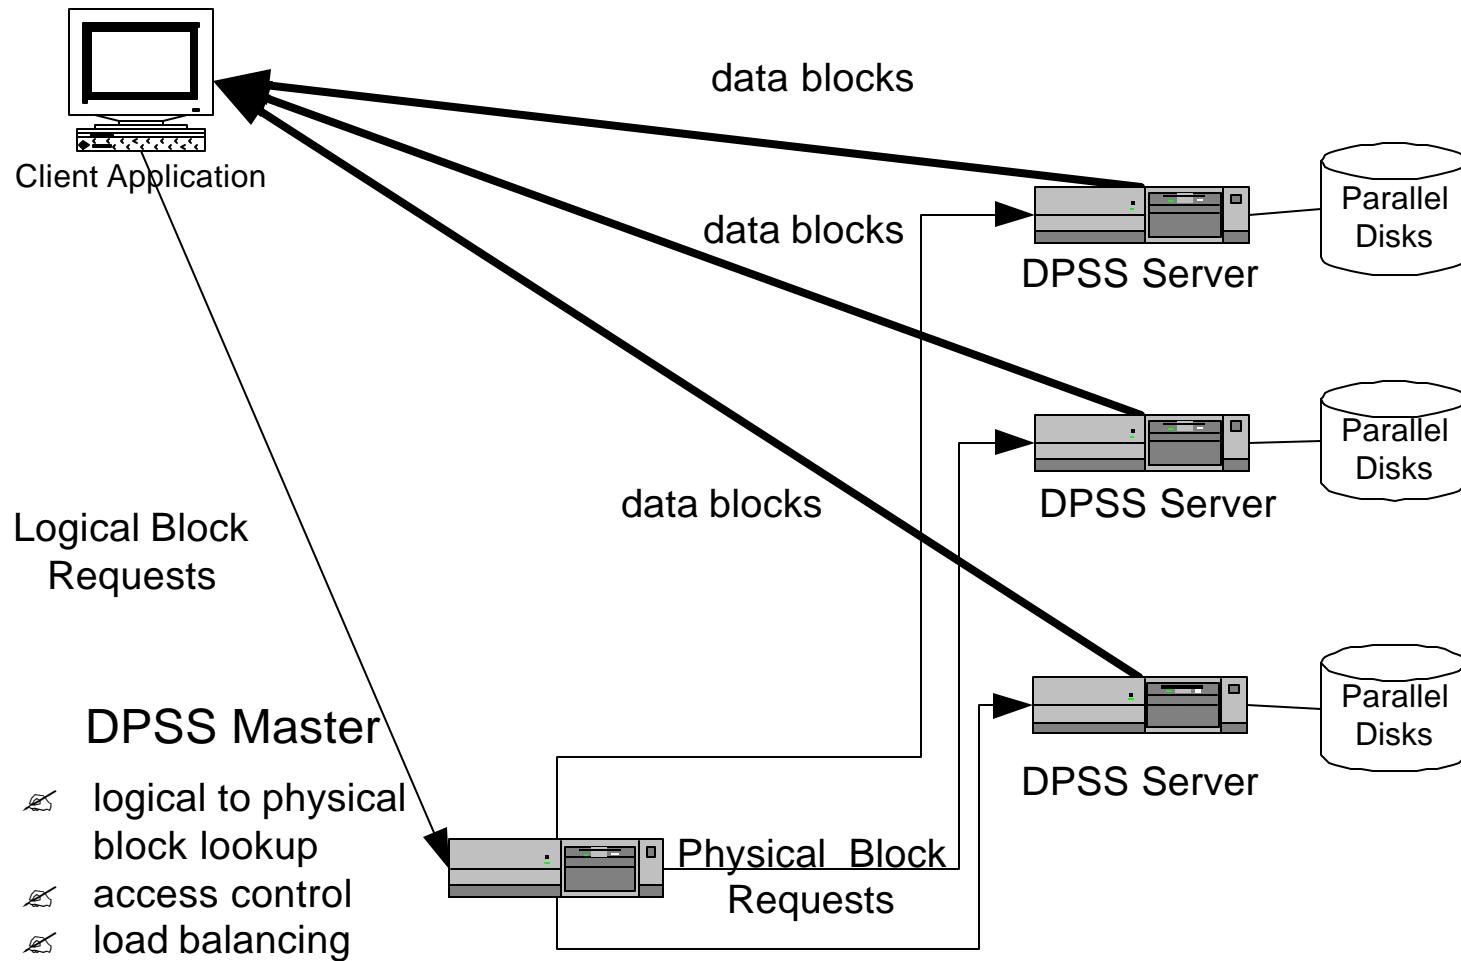
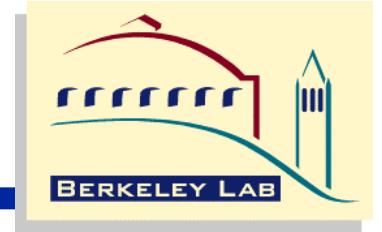
Other Results: LBNL to SNL-CA over NTON: sustained BW = 1.5 Gbps



Total Throughput (single dataset to a single cluster application):
1500 Mbits/sec (188 MB/sec) on 192 data streams, or 8 Mbits/stream



DPSS Cache Architecture

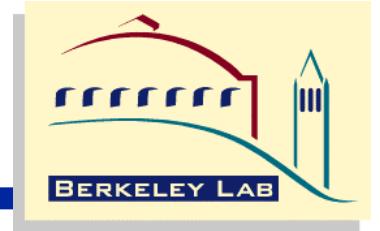


DPSS: 1 TeraByte High Performance Data Cache < \$12K (USD)



- Example: a 4 server DPSS system
- Sample DPSS Server Configuration:
 - Linux / Pentium (700 MHz PIII) server: \$1100
 - 3ware storage controller: \$300 (essential component!)
 - Netgear 1000BT card: \$310
 - 45 GB Western Digital IDE Disks; 6 per server: \$190/each
- Total Capacity: 24 disks x 45 GB = 1.08 TeraBytes
- Total Throughput: 250 Mbits x 4 servers = 1 Gbps
- Total Cost: ~\$2900 x 4 servers = \$11.6K

For more information



- DPSS info:
 - Brian Tierney: bltierney@lbl.gov
 - <http://www-didc.lbl.gov/>
- Visapult Info:
 - Wes Bethel: ewbethel@lbl.gov
 - <http://vis.lbl.gov/>